Properties of Retail Distribution and the Implications for Antitrust Analysis

Steven Tenn
Federal Trade Commission

and

John M. Yun*
Federal Trade Commission

November 16, 2004

* Contact information for John M. Yun (contact author): Federal Trade Commission, 600 Pennsylvania Avenue NW, Washington, DC 20580; phone: (202) 326-2433; fax: (202) 326-2625; and email: jyun@ftc.gov. We thank Chris Adams, Luke Froeb, Hajime Hadeishi, Sam Peltzman, and the participants of the EU/US Economic Workshop. This paper does not necessarily reflect the views of the Federal Trade Commission or any individual commissioner.

Properties of Retail Distribution
and the Implications for Antitrust Analysis

Abstract

Economic models of products sold in the retail channel typically only control for whether a product is available in at least one store in a given market, i.e., variation along the extensive margin. We demonstrate that it is also important to control for intensive margin variation regarding the depth of a product's availability. Looking at five grocery categories, we show that (a) many products are distributed in only a limited number of stores; (b) distribution undergoes significant variation over time; and (c) distribution's effect on sales is approximately constant returns to scale. These findings not only demonstrate that retail distribution is an important determinant of sales, but also displays significant variation that is not controlled for using only an extensive margin measure. We consider two applications where the importance of retail distribution has been under-appreciated: antitrust analysis and demand estimation.

Keywords: retail distribution, scanner data, demand estimation, portfolio effects, product cannibalization

JEL Classification: C81, D12, M3.

1 Introduction

Economic models of products sold in the retail channel typically only control for the extent of a product's retail availability in a binary manner. If a product is available in at least one store, then it is part of the vector of relevant products; otherwise, it is excluded from the analysis. This modeling choice recognizes that retail distribution is a necessary condition for retail sales. However, this "extensive margin" control for distribution does not consider intermediate or "intensive margin" variation in product availability. By solely focusing on whether a product is available at all, researchers ignore important variation in the depth of a product's availability. We find that intensive margin controls for distribution are virtually non-existent in the economic literature on the retail sector.

There are instances where an extensive margin measure sufficiently controls for a product's retail distribution. If a product is available in nearly all of the stores in a given market, or if a product's distribution is invariant over time, then the effect of product availability is absorbed into the intercept in many econometric specifications. In this case, omission of an intensive margin control would not affect the results.[1] Alternatively, one could start from the premise that distribution is relatively unimportant. This would be the case when search and transportation costs are sufficiently low that consumers make their purchase decisions based on the full set of products, even if individual retailers do not carry all of them.

The shared feature across these two possibilities is that the extent of a product's distribution has little power in explaining variation in retail sales, making its omission from the econometric specification unimportant. The empirical validity of this assumption has not yet been addressed in the literature. Looking at products in five grocery categories, we find there is little empirical support for not controlling for intensive margin variation in retail distribution.

[1] Tenn (2004a) demonstrates that when distribution is constant its omission does not affect own-price elasticity estimates in the constant elasticity and linear models of demand. However, estimates of cross-price elasticities in these models are biased when distribution is not controlled for.

We show that (a) distribution is significantly less than 100% for many products within a retail-chain and city; (b) distribution exhibits significant variation from period to period, even after controlling for seasonality and retailer-specific time trends; and (c) distribution's effect on sales is approximately constant returns to scale. Thus, not only is retail distribution an important determinant of sales, it displays significant variation that is not controlled for using only an extensive margin measure.

These results have important implications for economic models of the retail sector, and antitrust applications in particular. We comment on two recent cases. The first is the attempted baby food merger between Beech-Nut Nutrition Corporation and Heinz, which was successfully challenged by the U.S. Federal Trade Commission (FTC). The second is the acquisition of Seagram Spirits & Wines Group by Diageo and Pernod Ricard, which involved issues related to brand leveraging. Both cases demonstrate the potential benefits of incorporating distribution measures into empirical antitrust analyses.

We also consider the importance of retail distribution in demand estimation, which is widely employed not only in antitrust, but in many other areas of industrial organization as well. In particular, we undertake a cannibalization study that looks at whether changes in the distribution for a particular item affect the sales of a brand's other products. In a case study of cocktail franks, we find that greater sales from increased distribution primarily come from the "outside good," although there is a limited amount of same-brand cannibalization. The approach developed for cocktail franks is a general one that can be applied to other products as well.

The paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the dataset employed. Section 4 analyzes the empirical properties of distribution, including its level, variability, and effect on retail sales. Section 5 discusses the implications of the results, focusing on antitrust analysis and product cannibalization. Section 6 concludes.

2 Literature Review

The importance of product distribution is widely accepted in theoretical research concerning the retail sector. This is most readily apparent in the literature that analyzes slotting allowances. Slotting allowances are lump-sum payments that manufacturers make to retailers in order to gain distribution for their new products (FTC, 2003). This body of research recognizes that retailers have limited shelf space with which to stock products, and that consumers can only choose from among those goods that are available for purchase. These factors imply that product distribution is a major determinant of sales, and thus a valuable asset. The key slotting theories rely on these observations to explain why slotting payments from manufacturers to retailers might occur (e.g., Shaffer,1991; Sullivan, 1997).

Despite this appreciation that product distribution is a key determinant of retail sales, empirical investigations of the retail sector generally only control for distribution in a limited manner. Such an omission is unavoidable when datasets are employed which do not report detailed information concerning product distribution. However, the growing availability of retail scanner data allows progress to be made in this regard, since such datasets commonly include comprehensive information regarding product availability (see section 3).

When store and UPC-level data are employed, distribution is a binary variable; a given product is available at a particular store, or it is not.[2] Therefore, one can control for distribution by restricting the analysis to those UPCs available at a given store. Limited access to store-level scanner data, and the relative difficulty in undertaking UPC-level analysis, has led many researchers to conduct their investigations at more aggregate levels. In such studies, distribution potentially varies along both the extensive margin, whether a product is available in at least one store, as well as along the intensive margin, the number of stores that carry a given product.

[2] A "Universal Product Code," or UPC, represents a unique product, i.e., a particular combination of brand, flavor, package size, etc.

Consider, for example, Hausman and Leonard (2002), who study the competitive impact of a new brand introduction. The authors implicitly control for distribution in two distinct ways. They restrict their analysis to seven major brands of bath tissue that are (eventually) available in every city. Additionally, Hausman and Leonard separately analyze the data for the pre- and post-introduction periods. These specification choices allow the authors to maintain uniform product availability along the extensive margin in each part of their analysis.

However, Hausman and Leonard make no mention of whether product distribution underwent changes along the intensive margin. Not accounting for this margin of distribution is potentially problematic. The entry of a new product might not occur all at once, but rather through a gradual expansion across the stores in a given area. Further, only a few UPCs for a new brand might be initially introduced, with subsequent product-line expansions occurring over time.[3] Brand-level analysis that uses aggregate data obscures such changes in product availability, since they occur at a lower level of aggregation than the data employed.

Recent developments in discrete choice models of demand have led to an increasing focus on product-level analysis. Distribution along the extensive margin is easily controlled for in such studies, since discrete choice models allow consumer choice sets to differ across geographies or time (Ackerberg and Rysman, 2002).[4] Nonetheless, when such studies use data that has been aggregated across stores or UPCs, some margins of product distribution are still ignored. For example, studies often employ geographically aggregated datasets, such as city-level scanner data (e.g., Nevo, 2001). Other investigations aggregate different package sizes or flavors of a given type of item to form a single "product" (e.g., Shum, forthcoming; Villas-Boas, 2003). When not all stores carry every package size or flavor, distribution is left partially unaccounted for.

[3] In their research on product line extensions, Draganska and Jain (2003) find that the line lengths of various yogurt brands fluctuate and show "a fair amount of weekly variation" (p. 12).

[4] Looking at discrete choice models of demand, Ackerberg and Rysman fault the way studies control for product availability along the extensive margin. The authors propose a generalization that more flexibly allows the number of available products to enter the model specification.

That distributional changes along the intensive margin are less observable likely explains why such variation is commonly ignored. As such, the goal of this study is to make the importance of distribution more transparent, and thereby increase the likelihood that it will be accounted for in future analysis. To our knowledge, no previous study has systematically analyzed the properties of retail distribution. Empirical studies of distribution primarily come from the marketing literature. Farris et al. (1989), Reibstein and Farris (1995), and Bronnenberg et al. (2000) analyze the relationship between distribution and market share. Olver and Farris (1989) consider distribution in the context of push and pull marketing. Curhan (1972) indirectly looks at distribution in an analysis of shelf space's impact on unit sales. Draganska and Jain (2003) consider the impact of product assortment on the demand for yogurt. However, none of these papers analyzes the empirical properties of retail distribution in a systematic manner. In this paper, we try to fill this void by quantifying various aspects of distribution and by incorporating distribution into common retail applications such as demand estimation.

3 Data Description

This study utilizes supermarket scanner data provided by ACNielsen. The dataset contains weekly sales data for fourteen retailer-city combinations (e.g., the Jewel supermarket chain in Chicago) and comprises products in five grocery categories: hot dogs, frozen novelties, ice cream, shelf-stable pasta, and salad dressing.[5] The data is weekly from December 1998 to June 2001 (132 weeks).

The dataset contains dollar and unit sales for each UPC. In addition to weekly totals for each retailer-city combination, sales data is separately reported for four mutually exclusive levels of promotional activity: "No Promotion," "Feature Only," "Display Only," and "Feature & Display." A "Feature" is an advertisement, such as in a promotional circular, while a "Display"

[5] A confidentiality agreement with ACNielsen prohibits retailer names from being revealed. This example does not indicate whether the dataset employed contains the Jewel supermarket chain in Chicago.

is a secondary sales location within a store that is used to draw special attention to a given product.

The data reports two types of distribution: product distribution and promotional distribution. For each UPC, product distribution is measured using ACNielsen's "All Commodity Volume," or ACV, variable. ACV reports the percentage of total sales, across all product categories, accounted for by those stores that carry a given product. ACV represents the percentage of "standard-sized" stores that distribute a particular UPC. Similarly, the ACV for a given level of promotional activity is the fraction of sales, across all categories, accounted for by those stores where the UPC has said promotional activity.

While ACV is a measure of product-level distribution, some of our empirical investigations rely on higher levels of aggregation, such as brand-level analysis. To arrive at a distribution measure that corresponds to higher aggregation levels, we employ a variable known as "Total Distribution Points," or TDP. TDP is calculated as the sum of the ACV for each of the UPCs contained within a given set of products. At the brand-level, TDP is a combination of a brand's breadth of distribution (the percentage of stores that carry the brand) and the brand's depth of distribution (the number of distinct products carried per store). For example, a TDP of 300% is the distribution equivalent of a brand having three UPCs distributed at every store. At the subcategory and category level, TDP only captures variation in the depth of distribution, since essentially all supermarkets carry each of the product subcategories and categories contained in the dataset.

We consider three levels of aggregation in this study: UPC-level (or product-level), brand-level, and subcategory-level (e.g., super-premium ice cream is a subcategory of all ice cream). The use of multiple levels of aggregation allows for an analysis of both the breadth and depth of distribution. Further, as discussed in the following section, analysis at various levels of aggregation allows assessment of whether increased distribution for one product leads to the cannibalization of others, or whether it expands total brand or category sales.

4 Examining the Level, Variation, and Importance of Distribution

This section documents our key findings regarding various properties of retail distribution for the five grocery categories in our dataset. This analysis lays the foundation for the empirical applications considered in section 5. We find that most products are not carried by all of a retailer's stores in a given city. Rather, product distribution is well short of 100% for most UPCs. Similarly, but more surprisingly, products are often promoted in only a select number of stores within a retailer-city during a given week. That a particular item is on promotion in one store in a given retail-chain and city does not imply that all, or even most, of the other stores also have that product on promotion. In addition to this evidence that product availability and promotional activity are often limited, we show that product distribution frequently undergoes significant intertemporal changes, even after controlling for seasonality and retailer-specific time trends. Not only do individual products undertake significant gains and losses in distribution over short time periods, but also brands and even entire subcategories of products undergo major distributional changes. Lastly, we show that a product's distribution is a significant determinant of consumer demand. In fact, the evidence indicates that distribution is approximately constant returns to scale.

4.1 Distribution is Limited

Table 1 reports the product distribution level for each of the five grocery categories. Each observation represents a unique UPC that is available in at least one store (in a given week) for a particular retailer-city. For example, across the 249,192 observations in the frozen novelties category, the mean and median levels of product distribution are 70% and 80% of stores, respectively. The average distribution level across all five categories is 67% of stores, while the median is 76%. The results vary somewhat depending on the category. For pasta, 25% of the products are distributed in 18% of stores or less. In contrast, 50% of the products in the hot dog category have an ACV of 91% or more. Overall, the results indicate that while most products are carried by a sizable fraction of stores, relatively few items are available at all stores, and a significant fraction of products has very limited distribution.

Table 2 reports statistics regarding promotional distribution, where all three types of promotions are pooled together (i.e., "Feature Only," "Display Only," and "Feature & Display").[6] Each observation represents a UPC that was promoted in a given retailer-city for a particular week. For example, of the 17,468 promotional observations for frozen novelties in our sample, on average a product was promoted at 59% of the stores in a given retailer-city, with a median of 91% of stores. For all five categories combined, the average and median promotional distributions are 64% and 92%, respectively. In contrast to product distribution, many more products have a promotional distribution of 100%. In fact, the 75th percentile for all five categories is 100% of stores. However, despite it being common for a retailer to simultaneously promote a product in all of its stores in a given city, the results indicate that the assumption of universal promotions across a retailer's stores is inaccurate. It is not rare for a product to be promoted in only a small fraction of stores, such as for frozen novelties where a quarter of all promotions are in less than 10% of stores.

4.2 Distribution Exhibits Significant Variation Over Time

Since many products have limited distribution, this section investigates the extent retail distribution varies over time. To control for predictable changes in distribution, the following model is estimated:

(4.1) $\ln(TDP_{irt}) = \delta_{it} + \delta_{ir} + \alpha_{1ir}t + \alpha_{2ir}t^2 + \varepsilon_{irt}$, where $Var(\varepsilon_{irt}) = \sigma_i^2$.

For each product *i*, the log of TDP is regressed against a set of fixed effects for time *t* and retailer-city combination *r*, and a retailer-city specific quadratic time trend. The root mean

[6] This analysis could be performed separately for each type of promotion. However, by pooling the various types of promotions, any evidence that promotional distribution is less than 100% is made even stronger.

squared error (RMSE) from the regression, $\hat{\sigma}_i$, measures the extent distribution is changing over time after accounting for these factors.

A quadratic time trend is included in the specification to demonstrate that distributional changes are not solely reflective of gradual expansions or declines in product availability. Otherwise, the omission of distribution in analyses of retail sales data would not be problematic so long as one included time trend variables that accounted for such variation. Since quadratic time trends are commonly used (e.g., Chevalier et al., 2003), that is the specification we chose to employ.

Table 3 presents the RMSE from estimating (4.1) for three levels of aggregation.[7] For each aggregation level, the results are surprisingly similar across categories. The UPC and brand-level results show the most variation, each with an average and median RMSE of 28% across all five categories. The 25th percentile for the UPC and brand-level aggregations ranges between 13% and 27%, indicating that the vast majority of UPCs and brands undergo significant distributional changes over time.

The subcategory-level results reveal significantly less variation. The average RMSE is 12% across all five categories. This suggests that reallocation of shelf-space is more common within subcategories than across them.[8] Of the categories, ice cream and frozen novelties show the most variation, most likely because they have the largest number of subcategories. In unreported results, the RMSE at the category-level is 3%-4%, depending on the particular category. This indicates that shelf-space is even less fungible at the category level. While the space allocated to one type of frozen novelty (or ice cream) is sometimes given to a product in a different subcategory, the distribution of the entire category is almost entirely explained by the model's control variables.

[7] Equation (4.1) requires a sufficient number of observations to estimate the RMSE with reasonable precision. The approach taken here is to include all products with at least 156 weeks (three years) of observations (out of a maximum possible 1,848 pooled weeks). The results are not sensitive to this choice of sample selection.

[8] More precisely, since we include a quadratic time trend in the specification, the results indicate that changes in distribution are more "predictable" at the subcategory level rather than necessarily being less variable.

Thus, at the subcategory and category-levels of aggregation the retailer and time variables specified in equation (4.1) capture the vast majority of the variation in distribution. Therefore, empirical applications that operate at these higher levels of aggregation can successfully control for variation in product availability using such variables, rather than with an explicit measure of distribution. For levels of aggregation below the subcategory or category-level, however, the control variables in (4.1) are insufficient to capture a signification fraction of the variation in distribution. Rather, an explicit control for distribution is required.

4.3 Distribution Exhibits Constant Returns to Scale

The aim of this section is to determine the effect of increased distribution on sales. We start with a basic search model since the importance of product availability is directly related to consumer willingness to search for a specific item. To keep the analysis tractable, we consider consumer demand under two characterizations of search costs: costless search and infinitely costly search. The empirical analysis employs a generalization of these two models that contains each as a special case, thereby accommodating intermediate levels of search.

A complication in developing the model is the need to control for the impact of substitute products. Ideally, the demand specification for a given product should include the availability of all related products in all competing stores, as well as their price and promotional activity. Unfortunately, retailer-level scanner data does not report which competing products are available in each store. While the dataset contains information on the fraction of stores that carry each unique UPC, it does not report the joint distribution of each product's availability. Therefore, in the specification presented below, we only control for each product's own distribution, price, and promotional activity. Note, however, that we also estimate demand at higher levels of aggregation. At these levels, substitution across products within the same brand or subcategory is accounted for.

First, consider the case where search is infinitely costly. Consumers are limited to those products available at the store they typically frequent, or do not purchase anything at all.

Quantity sold q, for product i, in store s, at time t, is determined by a constant elasticity specification, where p denotes the product's price:

(4.2) $\ln q_{ist} = \mu_i + \alpha_i \ln p_{it}$.

As will become evident below, the advantage of employing a constant elasticity model is the natural way distribution can be incorporated into the specification. Summing demand across all stores within a retail-chain, which are assumed homogeneous, leads to the following specification, where π_{it} represents the fraction of stores that carry product i:

(4.3) $\ln q_{it} = \mu_i + \alpha_i \ln p_{it} + \ln \pi_{it}$.

Note that distribution is constant returns to scale, as evidenced by its unit coefficient.

Now consider the case of costless search. In this situation, consumers make their purchase decision using a choice set that includes all products available in any store. A direct implication is that a product's distribution has no impact on consumer demand (so long as it is available in at least one store). This assumption leads to the following demand model:

(4.4) $\ln q_{it} = \mu_i + \alpha_i \ln p_{it}$.

Models (4.3) and (4.4) differ solely with respect to the coefficient for the distribution variable $\ln \pi_{it}$. The following generalization incorporates these two models as special cases:

(4.5) $\ln q_{it} = \mu_i + \alpha_i \ln p_{it} + \gamma_i \ln \pi_{it}$.

If search costs are zero, and therefore distribution is unimportant, then $\ln \pi_{it}$ has a coefficient of zero. If search costs are infinite, then log distribution has a unit coefficient. Intermediate levels of search are represented by values of γ_i between zero and one.

Previous research demonstrates that aggregation bias often occurs when data is aggregated across stores with heterogeneous price and promotional activity (Christen et al. 1997; Link, 1995). We therefore modify model (4.5) by disaggregating the data by each type of promotional activity $m \in M$:

(4.6) $\ln q_{it}^{m} = \mu_{i}^{m} + \alpha_{i} \ln p_{it}^{m} + \gamma_{i} \ln \pi_{it}^{m}$.

As detailed in section 3, the set M contains four elements: "No Promotion," "Feature Only," "Display Only," and "Feature & Display." The variable π_{it}^{m} represents the fraction of stores where product i has promotional activity m, and is measured using the TDP variable.[9] Use of this specification avoids aggregation bias so long as a similar price is charged at all stores within a retailer-city that have the same promotional activity for a given product. This condition is substantially weaker than the price assumption often implicitly made when estimating demand using aggregate data, where prices are assumed identical across all stores regardless of their level of promotional activity.

Referencing (4.6), π_{it}^{m} can be rewritten as $\pi_{it}^{m} = \left(\pi_{it}^{m} \big/ \pi_{it}\right)\pi_{it}$, where $\pi_{it}^{m} \big/ \pi_{it}$ is equal to the fraction of stores that have promotion m for product i among those stores which carry that item. Thus, (4.6) can be rewritten as:

(4.7) $\ln q_{it}^{m} = \mu_{i}^{m} + \alpha_{i} \ln p_{it}^{m} + \gamma_{1i} \ln(\pi_{it}^{m} / \pi_{it}) + \gamma_{2i} \ln \pi_{it}$.

This equation allows for separate assessment of the impact of increased promotional distribution, γ_{1i}, and increased product distribution, γ_{2i}. The parameter γ_{1i} captures the impact of changing the mix of promotional activity across stores while holding the overall level of distribution constant. In contrast, γ_{2i} measures the impact of an increase in product distribution while holding constant the fraction of stores with each type of promotional activity.

Equation (4.7) is estimated for the five grocery categories and at the product, brand, and subcategory-levels of aggregation. To avoid aggregating across products that are too dissimilar, only items of the same package size are aggregated together.[10] Additionally, the model is estimated using time and retailer-city fixed effects and a retailer-city specific quadratic time

---

[9] Recall that at the UPC level of aggregation, TDP equals ACV.

[10] This avoids aggregation issues such as how to construct the average price across products that are measured in different units (e.g., ounces vs. the number of items per package).

trend. Also, similar to the previous subsection, the sample is restricted to those products where the model parameters can be estimated with reasonable precision.[11]

Table 4 presents the results. First consider the estimates for the promotional distribution elasticity γ_1. Strikingly, across the various categories and levels of aggregation, the average estimate across all products is approximately one. That is, promotional distribution is constant returns to scale. There is substantial heterogeneity in the point estimates across products, however. Depending on the category and level of aggregation, the standard deviation of the γ_1 estimates ranges from 13% to 42%.[12] Thus, while promotional distribution exhibits constant returns on average, many products display either increasing or decreasing returns to scale on an individual basis.

Table 4 also presents the estimates for product distribution elasticity γ_2. At the UPC and brand-levels of aggregation, product distribution exhibits slightly increasing returns to scale, on average. Like promotional distribution, the standard deviation of the estimates within each set of products can be quite high. At the subcategory-level of aggregation, the product distribution elasticity is closer to constant returns with an average estimate of 1.08 across all five product categories. In unreported results, further aggregation to the category level results in average estimates ranging from 1.03 to 1.10, depending on the category, none of which is statistically distinct from one at conventional levels. The lower elasticity estimates when using more aggregate data suggest that a portion of the sales expansion effect of distribution is due to the cannibalization of other products within a brand or subcategory. That is, a portion of the sales

---

[11] Again, we employ the criteria that each product has at least 156 weeks (three years) of data (out of 1,848 possible pooled weeks) and has elasticity estimates for promotional and product distribution with standard errors of no greater than one. Alternatively, we could use a weighted average to account for those observations that were imprecisely estimated. The weighted average is very similar to that obtained without using weights, but with the standard error criterion mentioned above. Only a small number of observations were dropped.

[12] When calculating the standard deviation of the estimates, we subtract the variation due to estimation error.

increase from the expansion of one item comes at the expense of related products' sales. This possibility is explored further in section 5.2.

In sum, not only does distribution change quite frequently over short periods of time, it has a significant impact on sales. This result is robust across categories, levels of aggregation, and both product and promotional distribution. The clear implication is that product availability is a key driver of retail sales. Further, given our results from sections 4.1 and 4.2, retail distribution is an important explanatory variable that cannot be controlled for using only an extensive margin measure. Rather, analyses of the retail sector must explicitly account for variation in product availability along the intensive margin.

Recall that in the context of the model we employ, a finding that distribution is important indicates significant consumer search or transportation costs. Previous research provides corroborating evidence that consumers find search quite costly. For example, Rhee and Bell (2002) use household panel data to undertake an empirical study of consumer mobility across supermarkets. They find a surprising lack of mobility, with consumers being very unlikely to change their primary store from one trip to the next. Given that the average household in their sample spends 94% of total expenditures at their main store for a given week, consumers must make the vast majority of their purchases over extended time periods at a single supermarket. Similarly, Bucklin and Lattin (1992) conclude that promotional activity does not persuade consumers to switch stores. Although Kumar and Leone (1988) do find evidence of inter-store substitution in response to temporary price and promotion changes, they conclude that geographic proximity is a significant factor that determines the magnitude of such effects. These findings imply a situation where consumers of grocery products face high search costs that impede their willingness to frequent other stores, even to purchase products that are on promotion. The evidence from Table 4, where promotional and product distribution are estimated to be close to one on average, similarly implies high search costs.[13] Thus, the findings presented in this section are not only consistent with previous research, but also provide new systematic evidence that controlling for distribution is important across a number of product categories.

5 Empirical Applications of Distribution

5.1 Antitrust Applications

In 2000, the U.S. Federal Trade Commission (FTC) challenged the acquisition of the Beech-Nut Nutrition Corporation ("Beech-Nut") by the H.J. Heinz Company ("Heinz"), both manufacturers of baby food products. This case demonstrates the potential benefits of incorporating retail distribution measures into empirical antitrust analyses. In particular, it illustrates a disjoint between the factual development of antitrust cases involving retail products, which often heavily rely on issues pertaining to retail distribution, and counterpart empirical analyses that incorporate retail distribution in only a limited manner.

In its complaint, the FTC alleged that a combined Heinz/Beech-Nut and the industry leader Gerber would together control 98% of the baby food market.[14] Despite the increased market consolidation, the district court found that a combined Heinz/Beech-Nut would enjoy significant efficiencies in production and would become a more effective competitor to Gerber.[15] Subsequently, the appeals court reversed the district court's decision citing insufficient evidence regarding these efficiencies.[16] The parties abandoned the proposed transaction soon thereafter.

---

[13] Table 2, which shows that products are only promoted in select stores, is also consistent with high search costs. If search costs were sufficiently small, then all consumers would purchase from the store(s) with the lowest price for a given item.

[14] *FTC v. H.J. Heinz Company and Milnot Holding Corp*., Memorandum In Support Of Plaintiff's Motion For Preliminary Injunction.

[15] *Federal Trade Commission v. H. J. Heinz Co.*, 116 F. Supp. 2d 190 (D.D.C. 2000), *rev'd* 345 U.S. App. D.C. 364, 246 F.3d 708 (D.C. Cir. 2001).

[16] *H. J. Heinz Co.*, 246 F.3d at 721.

In a review of the economic aspects of the case, Baker (2004) highlights the following critical fact: all three brands are rarely sold in the same store.[17] Competition for shelf-space was between Heinz and Beech-Nut to be the alternative to Gerber in a particular store. In an analysis of market-level sales data, Baker found insignificant or extremely low cross-price elasticities between Heinz and Beech-Nut. The district court judge explicitly relied upon this finding in his opinion.[18] In presenting these elasticities, Baker makes no mention of how the empirical specification integrated the fact that all three brands are rarely sold in the same store. From all appearances, Baker only controlled for distribution along the extensive margin, ignoring limited distribution along the intensive margin. This analytical shortcoming may explain why Baker estimated very little substitution between Heinz and Beech-Nut.[19]

Distribution played a large role in the factual and theoretical development of the case. However, it was largely ignored when estimating demand elasticities. In particular, Baker's empirical model implicitly assumes all stores within a particular metropolitan area are close substitutes; so long as a brand is available in at least one store in a given city, it does not matter whether it is available in all stores or only in a few of them. This assumption is contrary to our results, which strongly suggest that stores are heterogeneous and are frequented by consumers unwilling to expend a great deal of cost to search across stores.[20]

Another application where an intensive margin measure of distribution would contribute to antitrust analysis is Pernod Ricard S.A. and Diageo Plc's acquisition of Seagram Spirits & Wines Group in 2000. Seagram and Diageo are manufacturers of spirits such as rum and cognac.

---

[17] Baker served as the economic expert for Heinz and Beech-Nut.

[18] *H. J. Heinz Co.*, 116 F. Supp. 2d at 196.

[19] When intensive margin variation in distribution is ignored, it is impossible to distinguish between whether two products are poor substitutes when they are both available at the same store, or whether few retailers carry both of them (Tenn, 2004a).

[20] Thomadsen (2004) finds the same result for fast food outlets, which results in significant price dispersion within a given retail-chain across various locations.

In their deliberation on the competitive impact of the acquisition, the European Commission considered the possibility of brand leveraging, where "the ownership of a selection of leading brands may allow the brand owner to use secondary brands tactically *not as a revenue source* but also as a means of competing against other competitors' main brands" [emphasis added] (Case No. COMP/M.2268, Article 6(2), 08/05/2001, ¶23). The European Commission's concern was that a combined Diageo and Seagram would leverage their leading brands to garner more shelf-space for their secondary brands—at the expense of rival manufacturers. In order for this theory to be viable, two conditions must hold: (a) the acquisition must give the combined entity additional leverage it needs to successfully increase shelf-space for its secondary brands; and (b) this additional distribution must generate insufficient sales for retailers to have otherwise desired such an expansion. These conditions can be plausibly tested with distribution analysis. Post-merger data would reveal whether Diageo and Seagram's secondary brands increased their distribution at retail outlets, and whether this increase led to greater sales. Parallel to this analysis, one could analyze whether rivals' distribution decreased subsequent to the acquisition.

5.2 Product Cannibalization Case Study

Demand estimation is one of the most important types of analysis in industrial organization, playing a major role not only in antitrust, but in many other areas as well. The results of section 4.3, which show that retail distribution has a significant impact on sales, demonstrate the importance of incorporating an explicit measure of retail distribution when estimating the demand for a set of products. In this section, we focus on an important application of demand estimation: product cannibalization. In particular, we consider the effects of line extensions. For a more general treatment of how retail distribution can be incorporated into extant demand models, refer to Tenn (2004b).

A central issue a company must address when it is considering a line extension is whether introduction of new products will cannibalize the sales of their existing product line, or whether it will grow the brand's overall sales. In this section, we consider a case study of cocktail franks to obtain direct estimates of the degree of cannibalization—controlling for the distribution of each product within the line. This is an example of a horizontal line extension, where the new products have the same price and quality as the existing products.[21]

Cocktail franks are a niche product commonly used as a party appetizer. The advantage of analyzing this product is its simplicity; many supermarkets carry only a single brand of cocktail franks and generally carry only a handful of UPCs within that brand. We look at a subset of five retailers where one particular brand dominates, with all other brands having no or de minimis sales. Product distribution varies over time, with each retailer in a given city carrying between one and five unique UPCs for this brand. Each UPC differs with respect to the type of meat and flavoring used.[22]

For each UPC, we begin by employing a store-level constant elasticity demand model:

$$\text{(5.1)} \qquad \ln q_{ist}^{m} = \mu^{m} + \alpha \ln p_{it}^{m} + \delta_{c_{st}} .$$

The demand for each product i is a function of its promotion level m and price p. As before, there are four types of promotional activity (i.e., “No Promotion,” “Feature Only,” “Display Only,” and “Feature & Display”). The last term in equation (5.1), $\delta_{c_{st}}$, is a set of dummy variables that controls for the number of UPCs c that a given store s is currently carrying in week t. The products included in the dataset are line-priced.[23] Therefore, the price variable not only captures changes in a product’s own price, but also the price of the brand’s other UPCs.

The dataset does not report the multivariate distribution of UPCs across each of the stores in a retailer-city. To overcome this deficiency, we assume that the UPCs are independently

[21] In contrast a vertical line extension is primarily introduced to price discriminate according to consumers’ valuation of quality (Draganska and Jain, 2003).

[22] To enhance the comparability of the UPCs, we exclude from the dataset those products with nonstandard package sizes. These products represent a very small share of sales, making it unlikely that their exclusion has any meaningful impact on the results presented in this section. The UPCs included in the dataset all have the same package size.

[23] Meaning all of the UPCs for the brand are (nearly always) priced the same in the data.

distributed across all stores in a given retail-chain for a particular city and week. Under this independence assumption, we calculate the fraction of stores that carry a given set of UPCs using the ACV variable defined earlier as a measure of each product's distribution.

After summing equation (5.1) across all stores in a given retailer-city, we arrive at the following retailer-level demand model:

$$\text{(5.2)} \quad \ln q_{it}^{m} = \mu^{m} + \alpha \ln p_{it}^{m} + \ln \pi_{it}^{m} + \ln(\textstyle\sum_{c} \lambda_{cit} e^{\delta_c}) .$$

The variable λ_{cit} represents the fraction of stores, among those which carry UPC *i*, that stock exactly *c* UPCs, where $\sum_{c} \lambda_{cit} = 1$. The last term of (5.2) is the focus of the analysis, since it captures the cannibalization effect of the brand's other UPCs. This term is a function of the distribution of competing UPCs across the stores in a given retailer-city.

As before, the model implies that the coefficient for a product's own log distribution is equal to one. We generalize (5.2) to determine empirically whether this is the case. As in section 4, we differentiate between promotional distribution elasticity γ_1 and product distribution elasticity γ_2:

$$\text{(5.3)} \quad \ln q_{it}^{m} = \mu^{m} + \alpha \ln p_{it}^{m} + \gamma_1 \ln(\pi_{it}^{m} / \pi_{it}) + \gamma_2 \ln \pi_{it} + \ln(\textstyle\sum_{c} \lambda_{cit} e^{\delta_c}) .$$

Additional control variables employed are a set of UPC-retailer-city and week fixed effects, along with a retailer-city specific quadratic time trend. Equation (5.3) is a nonlinear function of the model parameters. We employ Quasi-Maximum Likelihood as the estimation method (Hamilton, 1994).[24]

Table 5 presents the results. All of the cannibalization parameters are negative, where the estimates represent the marginal effect relative to there being only a single UPC distributed at a particular store. These results indicate that the presence of additional UPCs of the same brand results in some cannibalization, i.e., negatively affects the sales of a given product. All but one

[24] That is, the first order conditions for maximizing the likelihood function, when the error term is i.i.d. normal, are used as moment conditions for GMM estimation. Newey-West standard errors are employed using a lag length of four weeks.

parameter are statistically significant at the 1% level. Importantly, at the 5% level of significance, one cannot reject that all four cannibalization parameters are identical. This indicates that almost all of the cannibalization is due to the first additional UPC.

The other estimated parameters are consistent with the results from section 4. The promotional and product elasticity estimates are large in magnitude and are statistically significant at the 1% level. At the 5% level of significance, one cannot reject that each estimate is equal to one, which is consistent with both types of distribution having constant returns to scale.

In summary, the evidence indicates that moving from one to two UPCs results in moderate cannibalization (approximately one quarter of unit sales). The impact of additional UPCs, however, is small or statistically insignificant. On net, the marginal impact of cannibalization is much less than the marginal impact of increased distribution. Therefore, increasing the distribution of a product within a certain brand increases the sales of the brand as a whole.[25] These findings are consistent with previous research that documents limited cannibalization from line extensions (Reddy et al. 1994 and, to a lesser extent, Lomax and McWilliam, 2001 who find mixed results). If the limited cannibalization that is found in this case study is indicative of other product categories, this might help explain why manufacturers often use line extensions rather than build new brands (Aaker, 1991).

Our results indicate that most of the additional sales from a line extension come at the expense of the "outside good," rather than from cannibalizing the sales of other products in the category. This is quite remarkable given the similarity of the products involved, which differ only by the type of meat and flavorings used. Our earlier finding that distribution is, on average, approximately constant returns to scale across various levels of aggregation, suggests that this may be an attribute of other supermarket products as well. Such a finding would have major

[25] Of course, from a retailer's perspective the relevant issue is whether increasing a product's distribution will maximize its profits, not whether doing so will increase total sales.

implications for antitrust analysis. When consumers primarily substitute to the “outside good,” even a product’s closest competitors may be relatively poor substitutes. Horizontal mergers involving such highly differentiated supermarket products would likely pose little antitrust concern.

6 Conclusion

This study is one of the first to consider systematically the characteristics of retail distribution. We find that distribution exhibits properties that are not widely recognized. Most products are distributed in far less than 100% of stores in a given retailer-city. Further, product availability varies significantly from week to week, and has an impact on sales that is approximately constant returns to scale. These features imply that incorporating an intensive margin measure of retail distribution could lead to significant improvements in applications such as demand estimation.

We find the widespread omission of product distribution in empirical analyses to be remarkable given our professional experience. Our impression is that, next to market share, “sales per distribution point” is one of the most common metrics brand managers follow. This likely explains why the major vendors of retail scanner data, ACNielsen and Information Resources Inc., report highly detailed information on product distribution. Our analysis of distribution did not require a “unique” dataset. Rather, the distribution variables employed in this study are widely available to the large number of academics and industry participants who utilize scanner data to analyze the retail sector (Bucklin and Gupta, 1999). Given the ease with which distribution can be incorporated into analyses of retail sales, we hope that a better understanding of the importance of distribution will lead researchers to pay greater attention to product availability.

While our study demonstrates several empirical regularities regarding retail distribution, it also raises a number of issues that still need to be addressed. In particular, what factors lead stores to carry different assortments of products? Why do retail-chains promote products in only some of their stores in a given week? We suspect that retailers make these choices in light of unobservable factors that influence demand, leading retail distribution to be an endogenous source of variation. Future research is required to determine whether valid instruments can be found to deal with this endogeneity, thereby extending the usefulness of integrating product availability measures into analyses of the retail sector.

References

Aaker, David A., *Managing Brand Equity* (New York: The Free Press, 1991).

Ackerberg, Daniel A. and Marc Rysman, "Unobserved Product Differentiation in Discrete Choice Models: Estimating Price Elasticities and Welfare Effects," Working paper, February 2002.

Baker, Jonathan B., "Efficiencies and High Concentration: Heinz Proposes to Acquire Beech-Nut (2001)," in *The Antitrust Revolution: Economics, Competition, and Policy*, edited by John E. Kwoka, Jr. and Lawrence J. White, 4^{th} edition, (New York, NY: Oxford University Press, 2004), pp. 150-69.

Bronnenberg, Bart J., Vijay Mahajan, and Wilfried R. Vanhonacker, "The Emergence of Market Structure in New Repeat-Purchase Categories: The Interplay of Market Share and Retailer Distribution," *Journal of Marketing Research* 37 (2000), pp. 16-31.

Bucklin, Randolph E., and Sunil Gupta, "Commercial Use of UPC Scanner Data: Industry and Academic Perspectives," *Marketing Science* 18 (1999), pp. 247-73.

Bucklin, Randolph E., and James M. Lattin, "A Model of Product Category Competition Among Grocery Retailers," *Journal of Retailing* 68 (1992), pp. 271-93.

Chevalier, Judith A., Anil K. Kashyap, and Peter E. Rossi, "Why Don't Prices Rise During Periods of Peak Demand? Evidence from Scanner Data," *American Economic Review* 93 (2003), pp. 15-37.

Christen, Markus, Sachin Gupta, John C. Porter, Richard Staelin, and Dick R. Wittink, "Using Market-Level Data to Understand Promotion Effects in a Nonlinear Model," *Journal of Marketing Research* 34 (1997), pp. 322-34.

Curhan, Ronald C., "The Relationship Between Shelf Space and Unit Sales in Supermarkets," *Journal of Marketing Research* 9 (1972), pp. 406-12.

Draganska, Michaela and Dipak Jain, "Product-Line Length as a Competitive Tool," Working paper, October 2003.

Farris, P., J. Olver, and C. De Kluyver, "The Relationship Between Distribution and Market Share," *Marketing Science* 2 (1989), pp. 107-28.

Federal Trade Commission, "Slotting Allowances in the Retail Grocery Industry: Selected Case Studies in Five Product Categories," FTC Staff Report, November 2003.

Hamilton, James D., *Time Series Analysis* (Princeton, N.J.: Princeton University Press, 1994).

Hausman, Jerry A. and Gregory K. Leonard, "The Competitive Effects of a New Product Introduction: A Case Study," *Journal of Industrial Economics* 50 (2002), pp. 237-63.

Kumar, V. and Robert P. Leone, "Measuring the Effect of Retail Store and Promotions on Brand and Store Substitution," *Journal of Marketing Research* 25 (1988), pp. 178-85.

Link, Ross, “Are Aggregate Scanner Data Models Biased?,” *Journal of Advertising Research* 5 (1995), pp. RC-8-12.

Lomax, Wendy, and Gil McWilliam, “Consumer Response to Line Extensions: Trial and Cannibalisation Effects,” *Journal of Marketing Management* 17 (2001), pp. 391-406.

Nevo, Aviv, “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica* 69 (2001), pp. 307-42.

Olver, James M. and Paul W. Farris, “Push and Pull: A One-Two Punch for Packaged Products,” *Sloan Management Review* 31 (1989), pp. 53-61.

Reddy, Srinivas K., Susan L. Holak, and Subodh Bhat, “To Extend or Not to Extend: Success Determinants of Line Extensions,” *Journal of Marketing Research* 31 (1994), pp. 243-62.

Reibstein, David J. and Paul W. Farris, “Market Share and Distribution: A Generalization, A Speculation, and Some Implications,” *Marketing Science* 14 (1995), pp. 190-202.

Rhee, Hongjai and David R. Bell, “The Inter-Store Mobility of Supermarket Shoppers,” *Journal of Retailing* 78 (2002), pp. 225-37.

Shaffer, Greg, “Slotting Allowances and Resale Price Maintenance: A Comparison of Facilitating Practices,” *RAND Journal of Economics* 22 (1991), pp. 120-35.

Shum, Matthew, “Does Advertising Overcome Brand Loyalty? Evidence from Breakfast Cereals,” *Journal of Economics and Management Strategy* (Forthcoming).

Sullivan, Mary W., “Slotting Allowances and the Market for New Products,” *Journal of Law and Economics* 40 (1997), pp. 461-93.

Tenn, Steven, “Demand Estimation Under Limited Product Availability,” Working paper, March 2004a.

Tenn, Steven, “Avoiding Aggregation Bias in Demand Estimation: A Multivariate Promotional Disaggregation Approach,” Working paper, November 2004b.

Thomadsen, Raphael, “The Impact of Location on Prices: Analysis in the Fast Food Industry,” Working paper, July 2004.

Villas-Boas, Sophia, “Vertical Contracts Between Manufacturers and Retailers: An Empirical Analysis,” Working paper, May 2003.

Table 1

Product Distribution

Category	N	% of Stores Mean	25th Pct	Median	75th Pct
Frozen Novelty	249,192	70%	53%	80%	94%
Pasta	434,558	56%	18%	60%	94%
Hot Dog	102,209	78%	68%	91%	100%
Ice Cream	515,303	66%	43%	73%	93%
Salad Dressing	494,591	63%	33%	74%	94%
Average		**67%**	**43%**	**76%**	**95%**

Notes: N = number of observations. Each observation corresponds to a UPC in a given retailer-city for a particular week. Sample restricted to observations with positive product distribution. Distribution is measured using the ACV variable defined in section 3.

Table 2

Promotional Distribution

Category	N	% of Stores on Promotion Mean	25th Pct	Median	75th Pct
Frozen Novelty	17,468	59%	10%	91%	100%
Pasta	26,732	59%	11%	82%	100%
Hot Dog	13,982	66%	13%	100%	100%
Ice Cream	76,219	74%	23%	100%	100%
Salad Dressing	39,659	60%	12%	87%	100%
Average		**64%**	**14%**	**92%**	**100%**

Notes: N = number of observations. Each observation corresponds to a UPC in a given retailer-city for a particular week. Sample restricted to observations with positive promotional distribution. Distribution is measured using the ACV variable defined in section 3.

Table 3

Root MSE of Log Distribution

UPC-Level Aggregation

	N	RMSE of Log Distribution				
		Mean	StdDev	25th Pct	Median	75th Pct
Frozen Novelty	332	27%	12%	17%	28%	37%
Pasta	818	27%	14%	15%	29%	38%
Hot Dog	151	25%	14%	13%	23%	34%
Ice Cream	679	29%	14%	18%	30%	40%
Salad Dressing	732	31%	13%	22%	32%	40%
Average		**28%**	**13%**	**17%**	**28%**	**38%**

Brand-Level Aggregation

	N	RMSE of Log Distribution				
		Mean	StdDev	25th Pct	Median	75th Pct
Frozen Novelty	100	28%	14%	18%	27%	35%
Pasta	157	28%	15%	18%	30%	39%
Hot Dog	46	26%	15%	15%	21%	38%
Ice Cream	75	27%	13%	17%	30%	35%
Salad Dressing	118	33%	13%	27%	34%	43%
Average		**28%**	**14%**	**19%**	**28%**	**38%**

Subcategory-Level Aggregation

	N	RMSE of Log Distribution				
		Mean	StdDev	25th Pct	Median	75th Pct
Frozen Novelty	12	22%	14%	10%	17%	35%
Pasta	4	5%	2%	3%	4%	6%
Hot Dog	2	11%	9%	4%	11%	17%
Ice Cream	15	20%	14%	8%	16%	26%
Salad Dressing	4	4%	1%	4%	4%	5%
Average		**12%**	**8%**	**6%**	**10%**	**18%**

Notes: N = number of observations. Each observation corresponds to a set of products at the level of aggregation specified. Distribution is measured using the TDP variable defined in section 3. The table reports the root mean square error from regressions using log TDP as the control variable (see text for details).

Table 4

Promotional and Product Distribution Elasticity Estimates

Category	N	Promotional Distribution Elasticity: Mean	StdDev	25th Pct	Median	75th Pct	Product Distribution Elasticity: Mean	StdDev	25th Pct	Median	75th Pct
UPC-Level Aggregation											
Frozen Novelty	269	0.97 (.02)	0.30	0.82	0.96	1.07	1.17 (.01)	0.19	1.04	1.17	1.27
Pasta	527	0.90 (.02)	0.39	0.73	0.94	1.09	1.16 (.02)	0.30	0.95	1.08	1.29
Hot Dog	133	1.02 (.03)	0.25	0.89	0.98	1.12	1.31 (.04)	0.39	1.07	1.26	1.43
Ice Cream	639	0.94 (.01)	0.16	0.86	0.96	1.06	1.28 (.02)	0.37	1.06	1.19	1.38
Salad Dressing	543	0.94 (.02)	0.34	0.85	0.96	1.05	1.19 (.02)	0.33	1.00	1.13	1.30
Average		**0.95 (.01)**	**0.29**	**0.83**	**0.96**	**1.08**	**1.22 (.01)**	**0.32**	**1.02**	**1.17**	**1.33**
Brand-Level Aggregation											
Frozen Novelty	126	1.01 (.03)	0.28	0.88	0.99	1.11	1.14 (.02)	0.19	1.01	1.14	1.25
Pasta	178	0.95 (.03)	0.34	0.83	0.95	1.10	1.16 (.04)	0.54	0.94	1.05	1.22
Hot Dog	62	0.99 (.04)	0.25	0.88	0.99	1.12	1.19 (.05)	0.32	1.00	1.17	1.37
Ice Cream	62	0.97 (.02)	0.13	0.89	1.00	1.06	1.37 (.07)	0.55	1.09	1.24	1.40
Salad Dressing	108	1.04 (.04)	0.39	0.92	0.99	1.10	1.12 (.04)	0.41	0.95	1.02	1.20
Average		**0.99 (.02)**	**0.28**	**0.88**	**0.98**	**1.09**	**1.19 (.02)**	**0.40**	**1.00**	**1.12**	**1.29**
Subcategory-Level Aggregation											
Frozen Novelty	48	1.10 (.04)	0.23	0.98	1.11	1.20	1.11 (.02)	0.15	1.02	1.14	1.20
Pasta	43	1.11 (.07)	0.42	0.94	1.05	1.19	1.00 (.05)	0.31	0.86	0.97	1.08
Hot Dog	15	1.09 (.05)	0.15	1.01	1.08	1.18	1.05 (.07)	0.24	0.86	1.07	1.30
Ice Cream	40	0.96 (.05)	0.22	0.85	1.05	1.13	1.20 (.04)	0.22	1.03	1.19	1.31
Salad Dressing	40	1.02 (.05)	0.24	0.88	0.99	1.15	1.02 (.04)	0.20	0.90	0.97	1.19
Average		**1.06 (.02)**	**0.25**	**0.93**	**1.06**	**1.17**	**1.08 (.02)**	**0.22**	**0.94**	**1.07**	**1.22**

Notes: N = number of observations. The standard error of the mean estimate is reported in parentheses. Each observation corresponds to a set of products at the level of aggregation specified. Only products of the same package size are aggregated together. Distribution is measured using the TDP variable defined in section 3. The table reports estimates from a constant elasticity demand specification (see text for details).

Table 5

Cocktail Franks, UPC-Level Demand Estimates

Variable	Est	SE	
Log Price	-1.14	0.28	***
Promotional Distribution Elasticity	1.11	0.07	***
Product Distribution Elasticity	1.27	0.16	***
Feature Only	-0.01	0.07	
Display Only	0.79	0.16	***
Feature & Display	0.40	0.27	
Cannibalization			
2 UPCs Available	-0.28	0.08	***
3 UPCs Available	-0.35	0.10	***
4 UPCs Available	-0.14	0.09	
5 UPCs Available	-0.32	0.11	***

Notes: N=2,331 observations. The RMSE from the estimation is .35. Significance levels are reported according to *=10%, **=5%, and ***=1%. A product's own distribution is measured using the ACV variable defined in section 3. The table reports parameter estimates from a constant elasticity demand specification (see text for details). The set of promotion dummy variables is measured relative to "No Promotion." The set of cannibalization dummy variables is measured relative to there only being a single UPC available in a given store.